Because Privacy: Defining and Legitimating Privacy in iOS Development

Katie Shilton and Daniel Greene
University of Maryland, College Park

Abstract
Privacy is a critical challenge for the mobile application ecosystem. US policy approaches to mobile data protection rely on privacy by design: encouraging developers to proactively implement privacy features to protect sensitive data. But we know little about how application developers define privacy, how they decide to implement privacy features, and what motivates them to consider privacy as a primary design value. This project investigates the discussion of privacy as a professional practice within a community of mobile application developers. Analyzing posts on an iOS forum reveals that privacy is a frequent topic of conversation in this community. This paper describes how iOS developers define and legitimate privacy, and reveals a challenge: iOS developers rely heavily on a definition of privacy provided by Apple which does not reflect current empirical or theoretical understandings of how users understand privacy. Understanding this challenge can help us shape better guidelines for privacy by design, and broach challenges to the widespread adoption of privacy by design principles.
Keywords: Privacy; mobile development; discourse analysis; values in design
doi: 10.9776/16229
Acknowledgements: This research supported by NSF awards CNS-1452854, SES-1449351, and a Google Faculty Research Award.
Contact: kshilton@umd.edu

1 Introduction

Mobile applications constitute a unique software development ecosystem. Mobile applications (apps, in the vernacular of the ecosystem) are easy to build and distribute, fostering a culture of independent software developers alongside larger companies. Mobile apps can collect a large variety of sensitive personal data. An ecosystem of apps, hardware, operating systems, and telecommunication companies collect increasing amounts of personal data from mobile app users. Mobile developers (devs, as they often abbreviate themselves) are frequently positioned as key decision-makers about consumer privacy, deciding what data to collect, and how to store and share that data.

Current US policy approaches to protecting this data rely on privacy by design: encouraging devs to proactively implement privacy features to protect sensitive data. But we know very little about how mobile devs talk about values such as privacy, how they decide to implement privacy features, and what personal or institutional factors might motivate them to consider privacy as a primary design value. This paper begins investigating these questions through a discourse analysis of an online forum on which iOS devs talk about their work. It uses this analysis to ask:

1. How do iOS devs construct privacy as a professional practice?
2. How is privacy legitimated as a design value in this community?
3. How is privacy delineated and defined in this community?

Answering these questions provides insights for scholars and regulators interested in ethics and policy self-governance. This paper helps us understand both devs' interest and knowledge of privacy practices, what structures that knowledge, as well as gaps that might be addressed by platforms or regulators.

2 Background

Though the devs we observed interacting on the iPhone Dev SDK forum largely worked independently, there was no question they shared a culture with its own norms and practices. They are part of a historically specific online community of what Kelty (2008) calls 'geeks': not a pejorative label but a group united around common practices of tinkering with communication technologies. Here, we are explicitly concerned with how that community engages with ethical issues, and how ethical issues connect to work practices of building and designing. The value of privacy, for example, is informed by a broader set of legal and cultural norms filtered through these devs' specific social position and practices.

2.1 Developers and Politics

For Kelty (2008), geekdom is a community that arose among power users of the early Web. Geek communities tend to prioritize the free and easy sharing of resources and advice. Despite new digital trappings, Kelty's research revealed that geeks hold traditional modernist values that sanctify progress and see skilled polymathy as the best way to reach it. Coleman (2012) builds on the work of Kelty and others to explore specifically political geek values. She finds a form of 'geek liberalism' that is as much

practice as belief: geeks emphasize free individuals as the center of the political universe, just like classical liberals, and pursue that freedom through code that functions as public speech. Free/Libre and Open Source Software (FLOSS) devs share code freely and, in the process, redefine that code not as private property but public speech. In turn, the practice of collective tinkering applied first to code is ported to law, where corporate privacy policies and national legal regimes are dissected and publicized in the manner of a new software update. Our project builds on this rich ethnographic tradition by showing how core liberal values like privacy are defined in a commercial development community—in contrast to the FLOSS communities on which Kelty, Coleman, and others focus.

2.2 Forums and Values in Design

Development communities are not only made up of their members and their speech or code. The places where they meet—in this case, an online forum—themselves affect who can say what, where, when and how. Platforms moderate spam, police or limit forum exchanges, and constrain the kinds of interactions which can exist (Gillespie, 2010). For example, in a classic pre-Web paper, Yates and Orlikowski (1993) found that while email exchanges of a distributed work group shared certain linguistic features with both oral and written communication, they also shared new features unique to the medium.

While technological forms structure devs' discussions, the reverse is also true—a core insight of the values-in-design literature since, at least, Winner's (1980) assertion that technological artifacts have politics baked into them. For example, Marwick (2013) shows how the core features of Web 2.0 social media—sharing updates that collapsed social contexts and lended themselves to quantification—were a product of the San Francisco social milieu and the privileged devs in it who felt they had nothing to hide and needed to relentlessly promote their products. The effects of baking values into technical products reverberate far outside the dev community. Introna and Nissenbaum (2000), for example, show how for-profit search engines shift the shape of the Web by prioritizing mainstream commercial sites. Tracing the process of when and how devs' values are coded into the technologies they build is difficult. Shilton (2013) developed the idea of *value levers* to theoretically guide this empirical question. Drawing on fieldwork with developers, Shilton (2013, 2015) showed how particular design practices, particularly those that encourage devs to focus on data flows and think and act like users, triggered devs' recognition of social values and encouraged them to see the design process as social, not only technical.

2.3 Conflicting Definitions of Privacy

Privacy frequently rises to the forefront of conversations about devs, consumers, and platforms (Urban, Hoofnagle, & Li, 2012). However, disagreements about how to define privacy impact all of these stakeholders. In the U.S., policy definitions of privacy have centered on Fair Information Practices: best practices for corporate data collectors that center on providing notice of data collection, choice for consumers to opt out, access to data upon request, data security, and redress of errors (Waldo, Lin, & Millett, 2007). These practices are meant to satisfy consumers' individual privacy preferences; privacy-sensitive consumers can (theoretically) opt out of data collection, or request to see their data. However, empirical research has documented the failure of notice and consent (Martin, 2013), and shown privacy to be less dependent upon individual preferences than social norms (Martin & Shilton, 2015). Privacy scholars increasingly argue that privacy requires understanding cultural norms dictating which information should be shared with whom, when and why (Dourish & Anderson, 2006). Nissenbaum has proposed an influential framework to define these norms: *contextual integrity*, or the governance of the flow of personal information through contextually-specific "norms, policies, law, and technical design[s]" (2004, p. 6). Contextual integrity focuses on *context-appropriate* information flows, acknowledging that people develop generalizable expectations about how information flows in specific contexts with time and experience (Nissenbaum, 2009). Those privacy norms dictate what information it is acceptable to collect, who can have access to it, whether it should be kept confidential, and how it can be shared and reused. Shopping online, playing a game, or divulging information to a doctor are each governed by different information norms. The importance of context to privacy has since transformed academic work and been incorporated into influential policy documents (NTIA, 2012).

Through their design decisions and interactions with Apple's designs and policy, iOS devs can support—or violate—contextual integrity for millions of users. This approach recognizes privacy not just as an individualistic claim to personal information but a series of social exchanges between diverse, often conflicting, parties. A social view of privacy frames privacy as a struggle for the management of social boundaries between business, the state, individuals, and communities (Cohen, 2012). Our work empiricizes these theoretical interventions, exploring how key actors find and redefine those social boundaries through work and community practices.

3 Method

To answer our research questions and trace the work of privacy construction, we have undertaken a critical discourse analysis of mobile developer forums. Critical discourse analysis is a qualitative method for analyzing the way that participants talk about their social practices (Leeuwen, 2008). Quoting Foucault, van Leeuwen writes: “Discourse is ... ‘a socially constructed knowledge of some social practice,’ developed in specific social contexts, and in ways appropriate to those contexts…” (2008, p. 6). Critical discourse analysis looks for the ways that written texts recontextualize social practice by representing social actors, action, time, space, legitimacy, and purpose.

The data we analyzed using this method was drawn from the iPhone Dev SDK forums, one of the most popular and widely-used iOS developer forums online. It features such topics as code sharing, programming tutorials, open discussion and marketing guidance. Unlike other Apple-related forums, it is meant primarily for devs and does not focus on device reviews or speculation on product announcement, instead focusing on development advice. Unlike the official Apple developer forum, it does not require an Apple-issued Developer Key to participate. This means participants appear to be more diverse than those in the official forum, in terms of experience and purpose for participating, and that non-dev participants (e.g., advertising network representatives searching for potential clients) sometimes intervene. In addition, devs indicate liking the forum’s feeling of independence from Apple.

We collected data based on the phenomenon of interest. We searched for threads containing the term 'privacy,' and analyzed those that included a discussion of privacy. We discarded threads where privacy was used as a keyword in an advertisement for an app, or instances where devs posted job ads and promised privacy for job applicants. We found 155 results in June 2015 (ranging from 2009-2015) that fit these criteria. We exported these forums to Dedoose as HTML files for coding.

Both authors performed a first read of the data to generate a set of initial thematic codes focused on privacy definitions. Right away we noticed legal definitions, platform-specific definitions, and ethical definitions of privacy. We then divided the data set in half and coded threads separately, reviewing each other's codes in weekly meetings to ensure mutual understanding and thematic coherence. During this process, the code set grew to include emphasis on the opposite of privacy (generally data collection and personalization needs), ways that privacy was authorized and legitimated, and conceptions of other actors in the ecosystem (Apple, service providers such as SDKs, and users).

Our university's IRB certified that the forum data gathered here qualified as public data and thus did not qualify for further IRB review. However, we believe that quoting participants violates the contextual integrity of the forum space; forum participants may not expect their posts to be used for research. To minimize this violation, we have altered participant handles and slightly altered quotations within this paper to reduce the searchability of specific exchanges. Alterations preserve the original meaning of posts, and all analyses were conducted on the original, unaltered quotations. We also plan to announce our ongoing work on the forum, offer a survey to participants to gather information on their professional backgrounds and histories with the forum, and draw on those responses for follow-up interviews.

4 Findings

Forum participants develop apps for Apple's iOS platform. They participate in the forum as individuals, asking questions and giving advice. While first-timers do exist and certainly post requests, the majority of participants seem keen to establish their existing professional identity, even if they remain pseudonymous. Devs signal experience by referencing apps they developed in the past, policies that have changed over time, larger companies for whom they work or whose software development kits (SDKs) they use to develop their own products, apps whose success they envied, and the evolution of popular coding languages or hardware platforms. Participation is open to anyone developing in the iOS ecosystem but, because this is a technical forum for exchanging resources and advising peers, a premium is placed on expertise. While there are shared jokes over notorious apps and a shared knowledge base, there do not appear to be many longstanding social ties. Devs do not reference past conversations, or refer novice users to specific expert participants. Forum participation seems most like an industry conference mixer: people who regularly work alone have a chance to compare notes.

Devs frequently reach out to each other for expertise in defining and understanding privacy. Privacy, like all topics within the forum, becomes tied to professional identity. Knowledge of privacy is a way to establish expertise, or admit the lack thereof. As one dev wrote:

> From what I understood, you can't access the camera roll like you want to, because it could be a possible privacy breach. Please correct me and make fun of me (if you so choose) if I'm wrong.

The most common types of discussions asked for help, provide directions, or provided advice. Devs also engaged in identity work, including discussing the state of their shared industry. Threads were often initiated by a participant who had run into a development hurdle and is seeking the advice of colleagues with similar experiences. These threads were initiated at multiple stages in the development process (code that won't run, advertising networks that aren't returning sufficient revenue), but in threads about privacy, seemed most concentrated at the point directly before or after an app is submitted to Apple for review. Threads were also initiated by participants early in the development process with general questions about what their app could or could not do (or, less frequently, should or should not do).

Forum exchanges were casual but curt, just-the-facts professionalism embellished with the aesthetics of geek culture: netspeak (e.g., lol, iirc, emoji and ASCII art), goofy avatars (video game characters, large banners reading "User Banned, Carry On"), disregard for capitalization and spelling, etc. Most exchanges were short: two or three sentences which respondents then quote or excerpt in their own response. Lengthier responses occasionally consisted of detailed, step-by-step technical advice. This practice was relatively rare, as the interface wasn't conducive to posting long strings of code. Most technical advice pointed to a specific trouble spot in the implementation process or an off-site resource. Lengthier responses frequently signaled thorny ethical or political issues to be worked through.

The forum is dominated by devs, but there are also outside actors whoinfluence discourse within the forum, whether by directly intervening or by their persistent presence in dev discussions. In privacy discussions, regulators and lawyers occasionally make an appearance. The main outside actors are Apple and its reviewers, and popular dev SDKs and their representatives in the forums.

A unique feature of the iOS ecosystem is Apple's role as a gatekeeper to the marketplace. Devs use the forum to figure out Apple's rules and practices to try to ease their path to market. "Apple" sometimes takes the form of individual app store reviewers, particularly when devs feel a rejection is unwarranted. More frequently, Apple as a corporation is the subject of discussion, as devs explain to each other that specific functions are enabled or disabled because of Apple's privacy rules. The values of the platform aren't so much described as discovered or delineated: lines are drawn and the consequences for crossing them elucidated, but the logic behind the line's placement is a secondary concern. The 2012 thread "How can we know if our feature crosses apple's border line", for example, is a back-and-forth between three devs about whether a feature is technically impossible or purposefully restricted based on the privacy features baked into iOS.

While Apple and its reviewers help define privacy within the forum without being present themselves, representatives from companies that sell services to devs (such as advertising platforms) actively participate in the forum. These outside actors add a new set of data sharing concerns, re-positioning the boundaries of privacy.

4.1 Why Privacy?

We first turned to analyzing why privacy was valued by forum participants, reviewing arguments that legitimated respect for, and design for, privacy. Building on categories identified by van Leeuwen (2008), we identified the telling of stories to illustrate good and bad consequences of ignoring privacy (mythopoesis); moral arguments for privacy (moral evaluation); technical and instrumental arguments for the importance of privacy (rationalization); and, most frequently, instances of appealing to the authority of governing bodies as well as users (authorization).

4.1.1 Mythopoesis

Van Leeuwen defines mythopoesis as: "…legitimation conveyed through narratives whose outcomes reward legitimate actions and punish nonlegitimate actions" (2008, p. 106). In dev forums, stories that legitimize privacy took two predominant forms: moral tales (which identify particular actors as bad), and cautionary tales (in which actors are punished).

A frequent moral tale was the invocation of "spyware". Spyware was always defined as immoral, and devs took pains to distinguish their apps from spyware. As a dev wrote when he asked if he could create an always-on location tracking app: "This will not be spyware, and the user will be fully aware of this feature if they launch the application." In a different thread focused on analytics tools, a user asked: "If you use either [SDK A] or [SDK B] are you supposed to tell your users? Its basically spyware in a way right?" A flurry of posts then sought to distinguish both SDKs from spyware.

Cautionary tales informed others of bad outcomes that could result from particular forms of data collection. Some cautionary tales imagined concrete (if unlikely) consequences for bad privacy decisions:

> Indeed, I think I would sue if I found out an app was filming me without my knowledge or permission. If you upload that video that would probably be felony invasion of privacy. (Read prison time.)

Other cautionary tales warned about bad actors. As a dev wrote in response to a poster who wanted to access phone call and text message logs: “Thank God these logs are inaccessible. That's a lot of privacy related data that shouldn't fall into the wrong hands.” Sometimes devs were more explicit in their imagining of bad actors. As one dev wrote in 2011:

> The truth? If someone wants to spy on you they will find a way. …BUT high-level spycraft takes work, your isp or someone else (the feds?? riaa??) would need to have some good reasons to waste time and resources on you.

In a 2011 thread, a European dev wrote:

> We are not located in USA or EU. We take privacy VERY seriously. I have denied to comply with subpoenas issued by US courts. None of the big companies in USA seem to do that. We have customers in the Middle East and other places to whom this is the main reason to choose [our application].This last point is something I have been struggling to get through, but the latest Wikileaks/Twitter subpoena case has given me some traction. It is safer to keep your data outside USA. People should and will take privacy more seriously in future.

Very direct cautionary tales such as this one, clearly related to Wikileaks revelations, were quite effective. The invocation of the US government as a privacy adversary prompted another dev to respond:

> That [privacy policy] should be the primary focus of [your] web page, in my opinion. …The title should be “We are the Swiss Bank of Email Providers." Seriously. People will get what that means in terms of their email security.

Another class of cautionary tales involved the reactions of users. As one dev wrote in a thread about metrics tracking apps: “What worries me is the possibility of user backlash against such tracking. There have been several famous cases where web tracking has got people in trouble.”

Other cautionary tales centered around privacy lawsuits, like this from a 2009 thread:

> This article is a warning for anyone that who do not play by the rule. From *PCworld*: “Lawsuit Claims IPhone Games Stole Phone Numbers”: "a pending class-action lawsuit filed against the devs, claiming that each of the company's games took advantage of a 'backdoor' method to access, collect, and transmit the wireless phone numbers of the iPhones on which its games are installed" ...The lawsuits are real and it will cost you a lot if you can not defend it or if you can not afford a lawyer. Let's begin the guessing game, how much "punitive damage" the lawyer want? 1 millions? 2 millions? May be declare bankruptcy before it finalized.

4.1.2 Moral evaluation

Some developers went beyond cautionary tales, which implied bad results for bad actors, and additionally made *moral* evaluations, in which invoking privacy was enough to shut down whole lines of development. As van Leeuwen describes it, moral evaluations represent:

> ...the tip of a submerged iceberg of moral values. They trigger a moral concept, but are detached from the system of interpretation from which they derive, at least on a conscious level… As a result, it is not possible to find an explicit, linguistically motivated method for identifying moral evaluations of this kind. As discourse analysts, we can only “recognize” them, on the basis of our commonsense cultural knowledge (2008, p. 110).

We coded tip-of-the-iceberg moral evaluations throughout the forums. Over and over again, devs told each other things like: “I don't think ur allowed to do that w/out prompting the user because of user privacy.” Privacy *was* the reason - it was enough all by itself, invoking moral concepts without having to go into the details of why and how. Invoking privacy could be enough to shut down a whole exchange:

> Dev1: Hi, I develop an app that needs to get the phone number of the device. So do you know the function that returns the iPhone phone number? Thanks
> Dev2: U cant do that. The privacy concerns associated for that would be insane
> Dev1: thanks

Developers frequently took strong moral stances. A dev in a thread about location data wrote:

> That's just impossible with the data from iTunes connect. The only way to do so should be sending you the device location at launch of the app but that would be against user privacy and therefore should not be done.

He later clarified: "That statement was my opinion not a policy related statement. I don't think it's right for devs to access that data if the app does not require it."

Sometimes devs took other participants to task for poor moral calculations. A dev upset with an advertising company for collecting what he deemed to be unnecessary address book information wrote a response to a representative:

> …sorry but I just do not buy this. U don't tell why u need the AddressBook framework and [you say there's] no way to have your platform without it. Yes I saw that also [a competing company] requires it. .... If your justification is that everyone does the same … It's like we steal cause many people also do steal.

Sometimes devs characterized other actors in clearly moral terms. For example, a dev wrote of a popular SDK: "You can go to their site ... and look at the video. sounds real 'evil'." Similarly, privacy was set up as a (rather extreme) political stance or belief that devs could hold. In a discussion about why a particular analytics tool was not (particularly) invasive, representatives of the tool clarified that the tool only collected aggregate user data. The tool representative qualified this, however: "As Coheed notes out, you may disagree with collecting any user data, which we respect."

4.1.3 Rationalization

Moral reasons were not the only arguments devs used to persuade others to care about privacy. Some devs blended moral evaluations with instrumental rationalizations. As one dev wrote:

> Personally, I'm against tracking every click/movement outside of beta testing. If [an app] was tested/designed properly then analytics at key points should be enough to pinpoint problems that turn up in production and determine how its being used.

Positioning themselves "against" full tracking was a moral stance, but it was backed up by an instrumental stance: such tracking was unnecessary if analytics were well-constructed.

Other devs rationalized privacy as a market necessity, believing users would abandon products violating user privacy. In a thread about how to capture video without users knowing, a dev wrote:

> Look, i havent looked into doing it, but based on what most people complain about i think that filming from their device without their knowing would be a big red light. It would for sure stop me from downloading an app if i saw/knew about that functionality.

The implication is that this sort of data collection behavior hurts sales.

4.1.4 Authorizing privacy

Although cautionary tales, moral authorizations, and rationalizations were common, developers most frequently legitimated privacy through *authorization* by a third party: by painting it as a requirement of governments, Apple, or users. References to government authority were rare, and often non-specific, as in a dev's declaration that: "The founding fathers ensured we had rights to personal privacy…" Occasionally guidance was more specific. As one dev wrote, referring to a quote from the Federal Trade Commission (FTC) in a popular press article:

> The FTC quote says it all: if you're developing mobile apps, you have to give the straight story about what your app can do and be transparent about your privacy practices.

References to Apple's authority were much more common. This was the single most-coded-for item: 151 statements (out of 2676 coded instances) were references to privacy authorization by the Apple platform. Examples frequently included privacy advice followed by a warning that noncompliance would trigger rejection: "You gotta ensure that the user knows and agrees to allow you to collect personal data or Apple may reject your app." In a discussion about automatically collecting email addresses, a dev put it memorably: "Wow, this is way worse than I thought. If Apple finds out what they did/are doing, they'll get into a **** storm." Devs displayed self-awareness about the potential for rejection. One responded to a thread: "Hmm lemme create a privacy policy fast or else my new update might get rejected."

Apple's privacy policies set design norms. In a discussion about getting rid of a location data awareness icon, a dev warned: "Apple wants the icon to appear, so the user knows that the GPS is running, both for power management and for privacy reasons. Getting rid of it is fighting Apple's design."

It was Apple's authority over privacy definitions that prompted frequent discussion and debate about where Apple draws privacy lines, and whether those lines could be negotiated. In a discussion of whether a dev should install an SDK that would track user analytics, one dev wrote: "I am very confident they'd begin rejecting Apps that have this sort of SDK in them." This confident statement prompted another dev to chime in and suggest discussing the plan with reviewers before beta-testing, to ensure their privacy standards were met before the project got further along.

Sometimes forum participants expressed that, not only did Apple authorize and necessitate privacy, but that they could *rely* on Apple to set privacy lines for them (and therefore justify any allowable data collection). For example, an ad network representative works to assuage dev ChillDude's privacy concerns about his network's SDK:

> ChillDude, Maybe I wasn't clear on what the address framework is used for: basically it…allows advertisers to serve ads that allow the user to interact with their address book. Think an ad that enables "email your friend" or "share with your family." But the user is always aware and opts in to these actions. Also, regarding apple, you can safely assume that if it's not against their policy, then you're pretty safe from violating any privacy concerns….This AddressBook linking has never given us an issue because it meets apple's seriously rigorous policies.

Although Apple's privacy guidelines were sometimes seen as protection, they could also serve as obstacles. A number of threads dealt with Apple's decision to deprecate the UDID identifier in 2011. The community was alerted to this change through a TechCrunch article posted to the forum. This discussion highlighted a challenge of Apple's dominant authority over privacy definitions: it was evidence that Apple's guidelines could change over time. Devs worried that new identifier solutions that didn't use the UDID would eventually face regulation. One dev offered: "I hashed the Mac address as an alternative, and it offers the exact same uniqueness (and so suffers the risk of being fought over in the future by Apple)."

Propensity for change was only one worry devs had about Apple's regulation. Many devs had trouble identifying when their designs would violate Apple policies, and recognized that Apple's policies left room for boundary negotiation. One offered this advice to a dev who had an app rejected for privacy concerns:

> Well according to the rejection letter you have to make [data collection] "clear to users". What that means to Apple I can't say with 100% certainty. But given industry privacy standards, as long as you note in your EULA stating that by using the app you're agreeing to the sending of the UDID [identifier], Apple should be appeased.

Devs thus acknowledged that Apple's privacy standards were not static. They worried that the line could shift over time, or based on which reviewer processed your app: an app approved today might be taken off the App Store next month, an App rejected in 2015 might be approved in 2016.

Changes were troubling to devs because Apple's authority was ultimate: devs felt a denied app was a waste of time and money. In a thread about rejections by Apple, user silvioc noted that keeping up with policy changes, and pivoting app designs based on those changes, was a serious resource drain for small companies. Other devs were resigned to this state of affairs: "It's like we're working with Apple. Apple is the boss. So when boss changed mind, then we should too."

Apple was not the only reference to authorize privacy. Users were also called upon as justification for privacy concerns. Sometimes user authority was expressed as an imperative to care for or watch out for users, while in other places privacy was invoked as a requirement put in place by users.

Protective measures stemming from an ethic of care for users were common. For example, a dev thinking of installing third-party advertising wrote:

> Hi, I was about to put [third party ad software] in an app I'm finishing, but I have a serious concern. AddressBook framework is required for having [company] ads. Excuse me? !!! I'm very concerned regarding my user's privacy. I want to know why this is needed for enabling [company] ads in my app. ... What data is collected? If Addressbook is needed I will not use [company software] ever.

Examples of care for users often invoked transparency or notice-and-consent requirements. Notice was an oft-repeated requirement: “If any personal information is being sent, the user should definitely be made aware (as is the case with location information).”

Sometimes users were invoked as an authority that would reject particular forms of data collection. In a thread about whether apps could harvest users’ email address, a dev replied:

> Users would run from you like a plague infested rat if they found out that you asked about this.... There is no need to steal their addresses & sell the list to the highest bidder!!Sorry for the harsh remark, but people think much too lightly on privacy today.

Particular audiences especially authorized attention to privacy. Unsurprisingly, apps for children had a strong privacy norm. But other audiences provided authority, too. As one dev wrote: “As well, since this will be used by lawyers and the like, protecting PII is really necessary.”

Debates over user privacy concerns also surfaced tensions between users and devs. In a thread about software that can help devs “watch how ppl use your app”, one dev wrote: “This seems good to devs, but, it obviously invades the privacy of users.” Users, however, are a complicated source of privacy authority. Some devs felt that notice and consent wouldn’t or couldn’t reach users:

> No one cares about privacy policies. They assume their data is safe already. Just saying "No personally identifiable information (PII) leaves the phone. This greatly reduces the risk of privacy problems" makes me wonder why I would have to worry about privacy problems in the first place. Why would privacy matter for a search app?

Interestingly, some devs felt that *Apple* assigned authority to users. In a discussion of an app rejection for privacy reasons, a dev wrote to the original poster: “That's have a valid point, but I doubt Apple will budge here, they care more about users than devs.” Often devs would draw on many ways to legitimize privacy in a single post. For example:

> Your question about Apple's policy is surely valid. However, regardless of what Apple says, it is worth having [a privacy] policy... Given they can be free and not really difficult to create, it just makes sense to do it. ...it’s a few sentences in a document. Maybe you collect [email address] for support, use it, and then delete it from storage. Just write that. It will look good to users and shows you’re serious about PII.

This dev acknowledges Apple’s authority on privacy matters, but also expresses that the dev’s moral duty extends beyond Apple’s requirements. He also expresses instrumental reasons to create a privacy policy (it’s not expensive or hard). Finally, he invokes user authority.

4.2 A Challenge: Defining Privacy

Forum discussions legitimizing privacy made it clear that privacy is an ethic for many developers. Privacy is something they value, respect, and develop for because they believe treating user data in particular ways is the right thing to do. However, what *constitutes* privacy is not so easily agreed upon. Developers spent a large amount of time trying to define and delineate what counts as privacy.

Maximalist definitions of privacy, in which any collection of personal data raised concerns, were surprisingly common. As one dev put it: “There is, still, the bigger ethical question of whether ANY transmission of ANY user data should be done without at least notifying the user.” Devs often defended themselves by denying collection of personal information: “We aren’t even using the location data…no personal data is sent back, we are just trying to determine use trends of the software.”

Also common are definitions of privacy which center on prohibitions against sharing data *beyond* the app company. Examples of these include promises that “...all data will be kept private.” These definitions also incorporate concerns about sharing personal data with third parties:

> My concern is, does Apple or other companies, like [SDK C] also collect user location including UDID [an identifier] at the same time when our own servers collect this info? If so, is it a privacy violation?

Another dev indicated that an app “sharing user location and data with third parties” was “on the edge of a privacy lawsuit.”

Transparency with users, particularly in the form of notice-and-consent, was the most frequently-used privacy definition. Devs frequently defined most kinds of data collection as allowable as long as users were informed. As dev Danny C described: “Are you asking if you can track the location of every

user for your app? The answer is yes, with the user's cooperation." Danny C continued with the stipulation: "You need a clear privacy policy that told the user that you were going to upload their location information, and let them know what you were going to do with that information.

5 Discussion

Privacy is a frequently discussed concern and a value of interest for devs in the iPhone Dev SDK dev forums. The many ways that devs define, authorize, and legitimate privacy illustrate a consensus over privacy as a concern, but little consensus over what 'counts' as privacy. Privacy is *valued* in the forums, but is only sometimes a *value* in terms of an articulated political philosophy. Privacy is more frequently a boundary object for practices and discourses around user data, structured by the app ecosystem—especially Apple's rules that prohibit apps with certain data collection practices from entering its marketplace. Veteran devs have more experience with those rules and are thus able to draw a better map of the boundary for novices. Certain data cannot be collected, shared, or sold *because privacy*.

But even as a negotiated limit, privacy has many functions within dev discourse. Privacy is discussed as an *explanation*. It can explain the opaque inner workings of Apple, or the preferences or behaviors of users. Privacy is also a *cudgel*. It prevents devs from doing everything they would like to, and punishes those who step out of line. Finally, privacy is an *ethic* for many devs. It is something they value, respect, and develop because treating user data in particular ways is the right thing to do.

Analysis of privacy discourse in the forums also revealed two important implications for technology ethics and policy. They illustrate that forums can be spaces of ethical deliberations. And it illustrates a "trickle-down privacy" effect in which platforms exercise strong power over privacy definitions.

5.1 Forums as a Space of Ethical Deliberation

Forum discussions illustrate that privacy inspires lively discussions, and that developers learn privacy norms from each other. Developers sometimes express surprise when data collections they'd planned struck others as creepy or invasive. As a dev considering a feature which would harvest users' email addresses wrote: "Thanks for the help. I realize I didn't consider the malicious way the info could be used." In a different conversation, a dev had a similar reaction: "Good points. I hadn't even considered the [invasive] image capture implications to this…" By serving as a space for ethical deliberation, the forum community can serve as a check on the first impulses of developers focused on data collection.

Findings from this study also have implications for ethics advocates who wish to intervene directly with technical communities (Brey, 2012; Shilton, 2015). Because privacy is a boundary object, it is also an important ethical object. If we want to talk to devs about ethics, privacy is a great place to start.

5.2 Trickle-down Privacy

IOS discussions about privacy illustrate that platforms can serve as an incredibly powerful regulator. It is clear that Apple's definitions of privacy (primarily focused on notice-and-consent) are widely adapted, mimicked, and quoted by devs in the Apple ecosystem. Devs were (accurately) much less worried about government regulation than regulation by this private corporation.

Unfortunately, there is a lack of transparency into how platforms regulate. Devs noticed this lack: the pervasive feeling that privacy was an always-moving target, and their energy spent trying to figure out the lines, reflect this lack of transparency. On the other hand, Apple can act more quickly than government, which in technology policy, can be an advantage. The power of Apple as a privacy regulator links to calls in information and media studies to understand the politics of platforms (Gillespie, 2010).

Privacy is of course *not* whatever Apple says it is, but by acting as a regulator, Apple narrows definitions of privacy. Apple policies are based on what many privacy scholars argue are outdated notions of notice and consent (Martin, 2013). Missing from both Apple's regulations and, importantly, dev conversations are discussion of alternate privacy models, such as privacy as contextual integrity.

5.3 Limitations and Future Work

Our exploration of privacy discourse in mobile development forums is just beginning. Future work will include comparison against other development ecosystems (such as the less-regulated Android platform), and investigation of the link between privacy discussions, work practices, and design decisions in mobile development. In addition, this work was limited by the keyword-search measures used to find privacy discussions in the expansive forums. Future work will consider how to find instances of privacy discussions that don't invoke the term "privacy." We also plan to compare forum conversations on other ethical values such as accessibility, fairness, and equity.

6 Conclusion

iOS forum participants debate and value privacy, and privacy is an important boundary object as devs negotiate the iOS ecosystem. At the same time, how to design for privacy remains contested within the space, and iOS devs' relative consensus around notice-and-choice definitions does not fit the nuanced definitions of privacy delineated in privacy scholarship. Apple's—and therefore devs'—focus on notice-and-consent is a good start. But this study illustrates that devs interested in privacy as a value, or privacy as authorized and defined by users, may benefit from access to more nuanced privacy by design frameworks. Regulators – whether government or platform – might consider tools and policies that promote contextual integrity and more nuanced definitions of privacy in mobile application development.

7 References:

Brey, P. A. E. (2012). Anticipatory Ethics for Emerging Technologies. *NanoEthics*, *6*(1), 1–13.

Cohen, J. E. (2012). *Configuring the Networked Self: Law, Code, and the Play of Everyday Practice*. New Haven & London: Yale University Press.

Coleman, E. G. (2012). *Coding Freedom: The Ethics and Aesthetics of Hacking*. Princeton, N.J. and Oxford: Princeton University Press.

.

Dourish, P., & Anderson, K. (2006). Collective Information Practice: Emploring Privacy and Security As Social and Cultural Phenomena. *Hum.-Comput. Interact.*, *21*(3), 319–342.

Gillespie, T. L. (2010). The politics of "platforms." *New Media & Society*, *12*(3).

Introna, L. D., & Nissenbaum, H. (2000). Shaping the Web: Why the Politics of Search Engines Matters. *The Information Society*, *16*(3), 169–185.Kelty, C. M. (2008). *Two bits: the cultural significance of free software*. Durham, NC: Duke University Press.

Leeuwen, T. van. (2008). *Discourse and Practice: New Tools for Critical Discourse Analysis* (1 edition). Oxford ; New York: Oxford University Press.Leon, P. G., Ur, B., Balebako, R., Cranor, L. F., Shay, R., & Wang, Y. (2011). *Why Johnny can't opt out: a usability evaluation of tools to limit online behavioral advertising* (No. CMU-CyLab-11-017). Pittsburgh, PA: Carnegie Mellon University.

Martin, K. E. (2013). Transaction costs, privacy, and trust: The laudable goals and ultimate failure of notice and choice to respect privacy online. *First Monday*, *18*(12).

Martin, K. E., & Shilton, K. (2015). Why experience matters to privacy: How context-based experience moderates consumer privacy expectations for mobile applications. *Journal of the Association for Information Science and Technology*.

Marwick, A. E. (2013). *Status Update: Celebrity, Publicity, and Branding in the Social Media Age*. Yale University Press.

National Telecommunications and Information Administration. (2012). *Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy.* Washington DC: The White House.

Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, *79*(1), 119–158.

Nissenbaum, H. (2009). *Privacy in context: technology, policy, and the integrity of social life*. Stanford, CA: Stanford Law Books..

Shilton, K. (2013). Values levers: Building ethics into design. *Science, Technology & Human Values*, *38*(3), 374 – 397.

Shilton, K. (2015). "That's not an architecture problem!": Techniques and challenges for practicing anticipatory technology ethics. In *Proceedings of the iConference 2015*. Newport Beach, CA: ACM.

Urban, J. M., Hoofnagle, C. J., & Li, S. (2012). *Mobile Phones and Privacy* (BCLT Research Paper Series). Berkeley, CA: University of California at Berkeley - Center for the Study of Law and Society. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2103405

Waldo, J., Lin, H. S., & Millett, L. I. (2007). *Engaging privacy and information technology in a digital age*. Washington, D.C.: The National Academies Press.

Winner, L. (1980). Do artifacts have politics? *Daedalus*, *109*(1), 121–136.

Yates, J., & Orlikowski, W. J. (1993). *Knee-jerk anti-LOOPism and other E-mail phenomena: Oral, written, and electronic patterns in computer-mediated communication* (MIT Sloan School Working Paper No. 3578-93). Cambridge, MA: MIT Sloan School of Management.